

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
TYLER DIVISION

MOSAID TECHS., INC. *et al.* §
§
Plaintiffs, § Civil Action No. 6:11-cv-173
v. §
§
FREESCALE SEMICONDUCTOR, INC. *et* §
al, §
§
Defendants. §

**MEMORANDUM OPINION AND ORDER CONSTRUING CLAIM TERMS OF
UNITED STATES PATENT NOS. 5,577,230; 5,724,505; 5,958,036; 6,141,762; 6,256,725;
7,051,306; 7,415,680; 7,945,885; AND 7,996,811**

This claim construction opinion construes the disputed claim terms in United States Patent Nos. 5,577,230; 5,724,505; 5,958,036; 6,141,762; 6,256,725; 7,051,306; 7,415,680; 7,945,885; and 7,996,811, as asserted in the above captioned case. A *Markman* hearing was held on October 16, 2012, to construe the disputed terms of the various patents. For the reasons stated herein, the court adopts the constructions set forth below.

I. CLAIM CONSTRUCTION PRINCIPLES

“It is a ‘bedrock principle’ of patent law that ‘the claims of a patent define the invention to which the patentee is entitled the right to exclude.’ *Phillips v. AWH Corp.*, 415 F.3d 1303, 1312 (Fed. Cir. 2005)(quoting *Innova/Pure Water, Inc. v. Safari Water Filtration Sys., Inc.*, 381 F.3d 111, 115 (Fed. Cir. 2004)). The court examines a patent’s intrinsic evidence to define the patented invention’s scope. *Id.* at 1313-1314; *Bell Atl. Network Servs., Inc. v. Covad Commc’ns Group, Inc.*, 262 F.3d 1258, 1267 (Fed. Cir. 2001). Intrinsic evidence includes the claims, the rest of the

specification and the prosecution history. *Phillips*, 415 F.3d at 1312-13; *Bell Atl. Network Servs.*, 262 F.3d at 1267. The court gives claim terms their ordinary and customary meaning as understood by one of ordinary skill in the art at the time of the invention. *Phillips*, 415 F.3d at 1312-13; *Alloc, Inc. v. Int'l Trade Comm'n*, 342 F.3d 1361, 1368 (Fed. Cir. 2003).

Claim language guides the court's construction of claim terms. *Phillips*, 145 F.3d at 1314. “[T]he context in which a term is used in the asserted claim can be highly instructive.” *Id.* Other claims, asserted and unasserted, can provide additional instruction because “terms are normally used consistently throughout the patent.” *Id.* Differences among claims, such as additional limitations in dependent claims, can provide further guidance. *Id.*

“[C]laims ‘must be read in view of the specification, of which they are a part.’” *Id.* (quoting *Markman v. Westview Instruments, Inc.*, 52 F.3d 967, 979 (Fed. Cir. 1995)). “[T]he specification ‘is always highly relevant to the claim construction analysis. Usually it is dispositive; it is the single best guide to the meaning of a disputed term.’” *Id.* (quoting *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1582 (Fed. Cir. 1996)); *Teleflex Inc. v. Ficosa N. Am. Corp.*, 299 F.3d 1313, 1325 (Fed. Cir 2002). In the specification, a patentee may define his own terms, give a claim term a different meaning than it would otherwise possess, or disclaim or disavow some claim scope. *Phillips*, 415 F.3d at 1316. Although the court generally presumes terms possess their ordinary meaning, this presumption can be overcome by statements of clear disclaimer. *See Sci Med Life Sys., Inc. v. Advanced Cardiovascular Sys., Inc.*, 242 F.3d 1337, 1343-44 (Fed. Cir. 2001). This presumption does not arise when the patentee acts as his own lexicographer. *See Irdet Access, Inc. v. EchoStar Satellite Corp.*, 383 F.3d 1295, 1301 (Fed. Cir. 2004).

The specification may also resolve ambiguous claim terms “where the ordinary and accustomed meaning of the words used in the claims lack sufficient clarity to permit the scope of the claim to be ascertained from the words alone.” *Teleflex, Inc.*, 299 F.3d at 1325. For example, “[a] claim interpretation that excludes a preferred embodiment from the scope of the claim ‘is rarely, if ever, correct.’” *Globetrotter Software, Inc. v. Elam Computer Group, Inc.*, 362 F.3d 1367, 1381 (Fed. Cir. 2004) (quoting *Vitronics Corp.*, 90 F.3d at 1583). But, “[a]lthough the specification may aid the court in interpreting the meaning of disputed language in the claims, particular embodiments and examples appearing in the specification will not generally be read into the claims.” *Constant v. Advanced Micro-Devices, Inc.*, 848 F.2d 1560, 1571 (Fed. Cir. 1988); see also *Phillips*, 415 F.3d at 1323.

The prosecution history is another tool to supply the proper context for claim construction because a patentee may define a term during the prosecution of the patent. *Home Diagnostics, Inc. v. LifeScan, Inc.*, 381 F.3d 1352, 1356 (Fed. Cir. 2004)(“As in the case of the specification, a patent applicant may define a term in prosecuting the patent.”). The well-established doctrine of prosecution disclaimer “preclud[es] patentees from recapturing through claim interpretation specific meanings disclaimed during prosecution.” *Omega Eng’g Inc. v. Raytek Corp.*, 334 F.3d 1314 (Fed. Cir. 2003). The prosecution history must show that the patentee clearly and unambiguously disclaimed or disavowed the proposed interpretation during prosecution to obtain claim allowance. *Middleton Inc. v. 3M Co.*, 164 F.3d 1372, 1378-79 (Fed. Cir. 1988)(quotations omitted). “As a basic principle of claim interpretation, prosecution disclaimer promotes the public notice function of the intrinsic evidence and protects the public’s reliance on definitive statements made during prosecution.” *Omega Eng’g, Inc.*, 334 F.3d at 1324.

Although, “less significant than the intrinsic record in determining the legally operative meaning of claim language, “the Court may rely on extrinsic evidence to “shed useful light on the relevant art.” *Phillips*, 415 F.3d at 1317 (quotation omitted). Technical dictionaries and treatises may help the court understand the underlying technology and the manner in which one skilled in the art might use claim terms, but such sources may also provide overly broad definitions or may not be indicative of how terms are used in the patent. *Id.* at 1318. Similarly, expert testimony may aid the court in determining the particular meaning of a term in the pertinent field, but “conclusory, unsupported assertions by experts as to the definition of a claim term are not useful.” *Id.* Generally, extrinsic evidence is “less reliable than the patent and its prosecution history in determining how to read claim terms.” *Id.*

Determining the claimed function and the corresponding structure of means-plus-function clauses are matters of claim construction. *WMS Gaming Inc., v. Int'l Game Tech.*, 184 F.3d 1339, 1347 (Fed. Cir. 1999). Claim construction of a means-plus-function limitation involves two steps. See *Medical Instrumentation and Diagnostics v. Elekta*, 344 F.3d 1205, 1210 (Fed. Cir. 2003). The court must first identify the particular claimed function, and then look to the specification and identify the corresponding structure for that function. *Id.* “Under this second step, ‘structure disclosed in the specification is corresponding structure only if the specification or prosecution history clearly links or associates that structure to the function recited in the claim.’” *Id.* (citations omitted). “While corresponding structure need not include all things necessary to enable the claimed invention to work, it must include all structure that actually performs the recited function.” *Default Proof Credit Card Sys., Inc. v. Home Depot U.S.A., Inc.*, 412 F.3d 1291, 1298 (Fed. Cir. 2005).

II. '230 PATENT TERMS

A. Overview of the '230 patent

The '230 patent is directed to a microprocessor architecture having two memory buses to interface two separate memory systems for instruction fetch and data read/write with respect to both memory systems. The architecture is illustrated in Figure 3:




FIG. 3

The program fetch unit 50 can read instructions from either memory 46 or 48. The data read/write unit 52 can read data from or write data to either memory also. In order to coordinate operations between units 50 and 52 and either of memory 46 or 48, an arbitration unit 54 is provided. The arbitration unit services the requests for memory accesses originating from either program fetch unit 50 or data read/write unit 52. The arbitration unit 54 is further diagramed in Figure 4:




FIG. 4

As shown, unit 54 includes a bus controller for each bus and a control process 56. As indicated, arbitration unit 54 controls which bus is used for a given instruction request or data access and also determines what actions to take when two requests to the same bus occur simultaneously. The control process 56 defines a "mapping" between the requests for memory access and the appropriate bus connections. Such "mapping" is exemplified by the truth table of Figure 5 as implemented in the flow diagram of Figure 6.

B. Disputed terms in the '230 patent: non-means-plus-function terms

1. "Simultaneously stored." Found in '230 patent, Claim 1

Plaintiffs argue that this term should be construed as "both contained," essentially arguing that the specification makes it clear that the advantage of the claimed invention is that a memory is capable of containing both program instructions and data. Thus, the term means that both program

instructions and data are both contained in a memory. Plaintiffs also argue that Defendants' construction is technologically impossible.

Defendants suggest that the term should be construed as "written at the same time." They frame the parties' dispute as being whether the term means an act of storing (as Defendants contend) or a state of being stored (as Plaintiffs contend). Defendants argue that Plaintiffs' proposal is redundant, insofar as Claim 1 first recites that both memory means are for storing program instruction and data. Thus, the term must mean something other than containing both program instructions and data. According to Defendants, every time the word "simultaneously" is used in the specification, it means that two actions occur at the same time. Defendants also argue that the prosecution history supports their construction, because the claims were amended to distinguish over the Yasui reference's arbitration of conflicts of simultaneous stores from actual simultaneous storing of data and instructions to a memory means.

The court agrees with Plaintiffs that the proper construction of this term, when read in view of the specification, is "both contained." Defendants' argument of redundancy lacks merit. Claim 1 specifies that both memory means have a capability to contain both program instructions and data, which means that the memory means can store but does not necessarily store both simultaneously. The claim language is of a scope that covers an apparatus wherein during operation a first memory means contains only data and not program instructions despite having a capability of containing both, and wherein the second memory means contains both data and program instructions simultaneously. The claim language "storing program instructions and data," read in view of the specification and in view of the "simultaneously stored" recitation, actually means storing program instructions and/or data."

Defendants' reference to use in the specification of "simultaneously" as meaning "at the same time" is in the context of a conventional Harvard architecture wherein during a bus cycle data memory "simultaneously" supplies data to the data read/write unit while the program memory supplies program instruction code to the program fetch unit. What Defendants cite to in the specification does not concern simultaneous stores to memory, but rather simultaneous requests to fetch. Moreover, the specification describes that when a request to fetch is made for both program instructions and data and there is no conflict (i.e., different buses are required), the requests can be executed simultaneously. But if a conflict exists (i.e., same bus is required), the control process selects a source based on pre-assigned priority. The specification does not support an operation wherein program instructions and data are written simultaneously over the same bus to a memory.

Defendants' characterization of the prosecution history is similarly misplaced. The Examiner's reasons for allowance, *see Doc. # 196-2*, first notes that Harvard architecture has two memories—each one exclusively storing instructions or data—and thus the prior art does not teach simultaneous storing of instructions and data into one memory. The Examiner found the distinction to be a basis to allow the claims over the prior art. With respect to Yasui, the Examiner identified that its disclosure was of a Harvard architecture (i.e., first and second memory means exclusive to either data or instructions) combined with a common memory that stores both data and instructions. The Examiner further identifies that conflicts in simultaneous stores to the common memory requires arbitration.¹ The Examiner then indicates that Yasui's inclusion of a common memory that stores both data and instructions would eliminate a need to simultaneously store data and instructions into a memory means of the Harvard architecture. These comments do not support Defendants' argument

¹ Yasui is a combination of Harvard architecture and Von Neumann architecture (shown in Figure 1 of the '230 patent).

that the Examiner was referring to data and instructions being written at the same time to one of Yasui's first and second memory means.

The applicant's amendment leading to the allowance, Doc. # 196-1, characterized Yasui as disclosing a RAM common space that could receive conflicting data and instruction fetches. Yasui was distinguished as failing to disclose or suggest a memory means for storing program instructions and data, and instead has two dedicated memories. The Examiner's comments do not track the applicant's remarks distinguishing Yasui very well, but what the applicant stated (not the Examiner's interpretation of what the applicant stated) is controlling.

The court therefore construes this term to mean "both contained."

2. "Request controller." Found in '230 patent, Claim 2

Plaintiffs suggest that this term should be construed as "controller that maps requests to the first and second interface busses," arguing that Claim 2 is directed to the implementation of arbitration unit 54 (Figure 4) and that the term identifies the controller that maps requests to the bus interfaces. Plaintiffs contend that Defendants' construction is wrong because it is concerns simply controlling the operation of the bus controllers.

Defendants propose that the term should be construed as "a control circuit located within an arbitration unit that controls the operation of each bus controller." Defendants characterize the dispute as being a description of the operation of the request controller, and point out that the specification gives a specific identification of the operation from which they draw their construction.

Claim 2 re-states the structure corresponding to the arbitrating means of Claim 1, and the term "request controller" is another identification of the control process 56 structure. Claim 2 does

further specify the operation of the control process to include a requirement for a “preassigned priority” capability.

The term “request controller” when read in view of the independent Claim 1’s arbitrating means is readily understandable. The “mapping” and “bus control” requirements in the parties’ constructions accrue from the claim language itself, particularly in view of the construction for the “arbitrating means” (discussed below). Construing the term further is unnecessary. *See, e.g., United States Surgical Corp. v. Ethicon, Inc.*, 103 F.3d 1554, 1568 (Fed. Cir. 1997) (“The *Markman* decisions do not hold that the trial judge must repeat or restate every claim term in order to comply with the ruling that claim construction is for the court. Claim construction . . . is not an obligatory exercise in redundancy.”).

The court therefore concludes that no construction is necessary.

3. “Program instruction fetch”/“Fetches of program instructions.” Found in ‘230 patent, Claims 1 and 2

Although initially disputed, the parties later agreed that no construction for this term is necessary. The court agrees, and also notes that Defendants’ initial proposal was incorrect because it imposed a further act of “loading” a fetched instruction, thereby improperly adding a limitation to the claim.

No construction of this term is necessary.

C. Disputed terms of the ‘230 patent: means-plus-function terms

1. “Means for arbitrating which of the first and second memory bus interfaces a particular program fetch or data transfer is to take place.” Found in ‘230 patent, Claim 1.

The parties agree that this is a means-plus-function term, and that the function is: arbitrating which of the first and second memory bus interfaces a particular program instruction fetch or data transfer is to take place. The parties dispute only the structure for this term:

Plaintiffs' Proposed Structure	Defendants' Proposed Structure
Arbitration unit 54 that includes control process logic circuitry that resolves conflicts using pre-assignable priorities.	The arbitration unit shown in Figure 4, including a control process represented by block 56 that implements the truth table in Figure 5 and the flowchart in Figure 6 and two bus controllers 58, 60. The control process 56 controls which bus is used for a given request for instruction or data access and defines a mapping between the requests originating only from the program fetch unit 50 or data read/write unit 52 to the appropriate bus connections.

Plaintiffs argue Defendants' identification of additional structure is overreaching because it includes structure that is not necessary to perform the specified function. Defendants argue that the specification clearly links the control process 56 to all actions necessary to "arbitrate."

Plaintiffs' construction is problematic, insofar as it adds functionality that is not recited in the claim language, i.e., resolves conflicts. Defendants are correct that the corresponding structure is identified by the detailed structure of the combined Figures 4, 5, and 6. However, Defendants' construction also provides additional description of operation that is not relevant in identification of structure (i.e., the second sentence of the proposed construction).

The court agrees that, properly truncated to identify only the corresponding structure, Defendants' proposal is correct. The structure for this term is: arbitration unit 54 in Figure 4 including control process 56 implementing the truth table of Figure 5 and the flowchart of Figure 6 and bus controllers 58, 60, and equivalents thereof.

D. Agreed constructions

The parties agreed on the following constructions:

1. "First memory means for storing program instructions and data." Found in '230 patent, Claim 1.

This is a means-plus-function term. The function is: storing program instructions and data.

The structure is: memory systems, including RAM, ROM, PROM, EPROM, FLASH, EEPROM, and RAM, and equivalents thereof².

2. "Second memory means, for storing program instructions and data." Found in '230 patent, Claim 1.

This is a means-plus-function term. The function is: storing program instructions and data.

The structure is: memory systems, including RAM, ROM, PROM, EPROM, FLASH, EEPROM, and RAM, and equivalents thereof.

3. "Means for requesting fetches of program instructions from the first and second memory means." Found in '230 patent, Claim 1.

This is a means-plus-function term. The function is: requesting fetches of program instructions from the first and second memory means. The structure is: Program Fetch Unit 50, and equivalents thereof.

4. "Means for requesting transfers of data between the processor and the first memory means and the processor and the second memory means." Found in '230 patent, Claim 1.

This is a means-plus-function term. The function is: requesting transfers of data between the processor and the first memory means and the processor and the second memory means. The structure is: Data read write unit 52, and equivalents thereof.

² In each of the parties' agreed structures for means-plus-function terms, they omit the phrase "and equivalents thereof." However, the court has added this phrase in for each structure, as it is required by the statute. 35 U.S.C. §112, ¶ 6.



5. “Instruction data.” Found in ‘230 patent, Claim 8.

This term means: “The output of self modifying code written to memory and subsequently executable by the processor.”

III. ‘505 PATENT TERMS

A. Overview of ‘505 patent

The ‘505 patent is directed to a processor capable of making a trace recording of program instructions that have been executed. The trace recording can be exported via a serial interface. An overall microprocessor is shown in Figure 1 and the hardware development system (HDS) block is further diagramed in Figure 2.



As shown, the microprocessor is interfaced to an external debug host by serial port 44. The HDS block enables continuous on-chip program tracing of the instructions executed by the processor core. A tracing record can be output to the debug host via the interface and port. The breakpoint processor in the HDS block monitors instructions received by the processor core in accordance with breakpoint instructions to stop program execution or trigger various processor functions.

B. Disputed terms in ‘505 patent: non-means-plus function terms

1. “Serial port.” Found in ‘505 patent, claims 1, 5.

Plaintiffs suggest that this term should be construed as “a port having limited bandwidth serial data outputs.” Plaintiffs argue that the specification contemplates that the program trace can be adapted to various processor devices having limited bandwidth serial data outputs. Plaintiffs point to prior art wherein data is serialized from a single multi-bit block transmitted over a single conductor into multiple blocks of smaller bit pieces (e.g., three bits) and transmitted over three conductors. According to Plaintiffs, what is indicated is that the serial interface is a port having limited bandwidth serial data outputs.

Defendants argue that the term should be construed to mean “a port that transfers data over a single conductor one bit at a time.” Defendants suggest that the term has a standard meaning and Plaintiffs seek to deviate from it. According to Defendants, Plaintiffs’ construction effectively rewrites the claim to more broadly read plural “data outputs.” Defendants contend that the scope of the specification in identifying a “JTAG port,” which indicates, like all serial ports, data transfer is over a single conductor one bit at a time.

The court agrees with Defendants. Plaintiffs’ construction is not crafted from a view of the term in the context of the ‘505 patent specification. The term is one of art, which, as Defendants say, means a single conductor with data being sent one bit at a time. A serial port is in contrast to a parallel port. In a parallel port, multiple pins are used to send several bits at the same time. For example, an 8-bit byte can be send over 8 conductors with each conductor carrying a single bit of the byte. A serial port sends each bit of the byte one bit at a time in sequence, i.e., serially.

The court construes this term to mean “a port that transfers data over a single conductor one bit at a time.”

2. “Return address.” Found in ‘505 patent, Claim 1.

Plaintiffs contend that no construction is necessary, because the claim language itself adequately describes the meaning of the term. If the court does construe the term, Plaintiffs propose “the address to which a discontinuity instruction returns after being processed.”

Defendants suggest “the address of the program instruction from which a discontinuity event originated,” arguing that this comes directly from the specification.

The parties’ dispute centers around whether the term designates the address of a program instruction that produced a discontinuity event (Defendants) or designates an address where the program returns after a discontinuity instruction is processed (Plaintiffs). The specification describes recording the address of an instruction which originated a discontinuity event and labels that address as (RETURN_ADDR). ‘505 patent at 5:52-55. Defendants’ proposed construction is therefore consistent with the specification.

The court construes this term as “the address of the program instruction from which a discontinuity event originated.”

3. “Compressed program trace”/“compressing said program trace.” Found in ‘505 patent, Claim 1.

Plaintiffs argue that these terms do not need to be separately construed, given that they are part of the “receiving means” term that will be separately construed below.

Defendants suggest that the term be construed to mean “the minimum details about the discontinuity and conditionally executed instructions and their addresses necessary for the user to reconstruct a full program trace which can include the destination address, return address and instruction type being encoded to reduce its size. They argue that this comes from the specification, and a construction is needed to assist the jury.

The court disagrees with Defendants. Their proposal rewrites the claim language, and the terms “compress” and “compressing” have a commonly understood meaning (“to reduce in size”) which will be readily understandable to a jury.

No construction of this term is needed.

C. Disputed terms in the ‘505 patent: means-plus-function terms

1. “First receiving means . . . for receiving a first signal indicative of the one of said plural types of instructions processed by the processor core.” Found in ‘505 patent, Claim 1.

The parties agree that this is a means-plus-function term, but disagree as to both function and structure:

Plaintiffs’ Proposal	Defendants’ Proposal
<p>Function: receiving a first signal indicative of one of said plural types of instructions processed by the processor core.</p> <p>Structure: Line 30 (as depicted in Figures 1 and 2) and TBC block 50.</p>	<p>Function: receiving a signal indicative of one of a plurality of type of instructions processed by the processor core.</p> <p>Structure: Instruction type line 30, TBC block 50 in conjunction with INSTR_TYPE FIFO 54.</p>

According to Plaintiffs, the specification clearly links line 30 to the specified function. Plaintiffs criticize Defendants’ construction as being overly inclusive of structure that is not necessary to perform the function. Defendants point out that line 30 is described as transmitting data and not receiving data. Defendants argue that the specification makes clear that TBC block 50 receives the data.

As Defendants point out, the portion of the specification cited by Plaintiffs actually says that line 30 enables the processor core to transmit to the HDS block 26. Thus, contrary to Plaintiffs’ contention, the specification does not clearly link the function to line 30. Signals indicative of the types of program instructions executed by the processor core are conducted over line 30 but they are

received at trace buffer control (TBC) block 50 as Defendants contend. However, contrary to Defendants' argument to include FIFO 54, the specification describes its function to be temporarily storing instruction type data.

The court concludes that the function is as specified in the claim and proposed by Plaintiffs: receiving a first signal indicative or one of said plural types of instructions processed by the processor core. The structure is: TBC block 50 and equivalents thereof.

2. "Trace recording means . . . for conducting a program trace of the processor core."
Found in '505 patent, Claim 1.

The parties agree that this is a means-plus-function term, but disagree as to both function and structure:

Plaintiffs' Proposal	Defendants' Proposal
<p>Function: conducting a program trace of the processor core.</p> <p>Structure: trace buffer control 50, address FIFO 52, instruction type FIFO 54, and multiplexer 56 (as depicted in Figure 2).</p>	<p>Function: operable for conducting a program trace of the processor core.</p> <p>Structure: Arbitrator block 22, and the "loading stage of HDS block" described in Figures 3-5 and 6:40-8:42 including "OVERFLOW indicator stored in the INSTR_TYPE FIFO 54", flags "AF_FULL, ITF_FULL, FIFO_EMPTY, trace recording hardware 26, address FIFO 52, INSTR_TYPE FIFO 54, trace buffer control 50.</p>

Plaintiffs argue that the specification links the structure they identify to the program trace conducting function, while Defendants contend that the structure they identify is linked in the specification to the function.

The arbitrator block 22 provides data and instructions to the processor core 12. Thus, arbitrator block 22 is not linked to conducting a program trace. As the parties agree, the FIFO 52 and FIFO 54 are linked to program tracing. But, contrary to Defendants' contention, the remaining

structure identified in their proposed construction concerns a function of controlling the FIFO's operations for data throughput via the interface 24. '505 patent at 6:46-7:9.

The court concludes that the function is as specified in the claim and proposed by Plaintiffs: conducting a program trace of the processor core. The structure is: FIFO 52 and FIFO 54, and equivalents thereof.

3. "Trace compression means being operable to compress said program trace." Found in '505 patent, Claim 1.

The parties agree that this is a means-plus-function term, and that the function is: compress[ing] said program trace. The parties disagree as to structure:

Plaintiffs' Proposed Structure	Defendants' Proposed Structure
Trace buffer control 50 (as depicted in Figure 2).	Address FIFO 52, INSTR_TYPE FIFO 54, the portion of TBC block 50 which reduces the size of the INSTR_TYPE by implementing at least a Huffman encoder.

According to Plaintiffs, the specification describes compression as being an act to discard, rather than record, addresses and instruction types in order to reduce the information necessary for a trace. Plaintiffs criticize Defendants' identification of structure as importing limitations of the preferred embodiment.

Defendants point out that Claim 2 confirms inclusion of FIFO 52 and FIFO 54. Further, they argue that the only compression structure is the Huffman encoder, and that Plaintiffs' argument against including the Huffman encoder because it imports a limitation of the preferred embodiment is without merit in the context of a means-plus-function limitation.

Contrary to Plaintiffs' contention, discarding rather than recording is not linked to compressing the program trace. The program trace that is compressed is necessarily the recorded addresses and instruction types. The specification describes discarding an address as being in relation

to discontinuities and conditionally executed instructions, which may not be useful in a program trace. '505 patent at 5:31-6:9. As Defendants identify, compression of the recorded program trace is described in the specification as being provided by Huffman encoding applied by TCB 50 "to minimize the size of the signal which will eventually be transmitted by the JTAG interface 24.'505 patent at 6:12-27.

The court therefore concludes that the function of this term is: compress[ing] said program trace. The structure is: TCB 50 providing Huffman encoding, and equivalents thereof.

4. "Output means for sending." Found in '505 patent, Claim 1.

The parties agree that this is a means-plus-function term, and essentially agree to the function: Plaintiffs propose "sending out said compressed program trace via the serial port" and Defendants propose "sending out the compressed program trace via the serial port." Plaintiffs' proposal tracks the claim language—"said" rather than "the" compressed program. The function is therefore: sending out said compressed program trace via the serial port.

The parties disagree as to structure:

Plaintiffs' Proposed Structure	Defendants' Proposed Structure
Interface 24 to serial port 44 (as depicted in Figure 1).	Address FIFO 52, INSTR_TYPE FIFO 54, Multiplexer 56, line 38, the "output stage" portion of TBC Block 50 implementing Figures 6-7 (as described at 8:43-9:46) including TRACE_CAPTURE signal, the flags AF_FULL, ITF_FULL, and FIFO_EMPTY.

Plaintiffs contend that the structure linked to the function is JTAG interface 24, while Defendants argue that their identified structure is clearly linked to the function of sending data out the serial port.

Defendants point to the operation of the output stage of the TBC block 50. However, contrary to Defendants' argument, nowhere does the specification link that operation to the function of sending out a compressed program trace via the port 44. Instead, the specification describes that the on-chip program trace is transmitted to the debug host via the JTAG interface 24 and JTAG port 44. '505 patent at 5:17-19.

The court therefore concludes that Plaintiffs' proposal is correct. The structure is: interface 24 to serial port 44 (as depicted in Figure 1), and equivalents thereof.

5. "First control means . . . for controlling the operation of said trace recording means . . . being operable for asserting a start signal and a stop signal to said trace recording means such that when the start signal is asserted, said program trace is triggered by said first control means by enabling said trace recording means to receive instruction types from said first receiving means and addresses from said second receiving means and, when end signal is asserted, said program trace is terminated by deactivating said trace recording means." Found in '505 patent, Claim 7.

The parties agree that this is a means-plus-function term, and that the function is: control[ling] the operation of said trace recording means . . . being operable for asserting a start signal and a stop signal to said trace recording means such that when the start signal is asserted, said program trace is triggered by said first control means by enabling said trace recording means to receive instruction types from said first receiving means and addresses from said second receiving means and, when end signal is asserted, said program trace is terminated by deactivating said trace recording means.

The parties disagree as to the structure:

Plaintiffs' Proposed Structure	Defendants' Proposed Structure
Breakpoint processor 58 (as depicted in Figure 2).	Debug host computer 100 in conjunction with JTAG port 44 and JTAG interface 24; breakpoint block 58.

Plaintiffs argue that the specification clearly links breakpoint processor 58 to the stated function, while Defendants argue that the specification also links the debug computer 100 to the stated function.

Defendants are incorrect, as their proposal would improperly add additional structure. Plaintiffs are correct. The structure for this term is: breakpoint processor 58 as shown in Figure 2, and equivalents thereof.

D. Agreed constructions

The parties have agreed as to the following constructions:

1. “Receiving . . . a first signal indicative of the one of said plural types of instructions processed by the processor core.” Found in ‘505 patent, Claim 1.

The parties agree that this term carries its plain and ordinary meaning.


2. “Second receiving means . . . for receiving a second signal indicative of a destination address of an instruction processed by the processor core and, when said instruction is of a discontinuity type, a return address of said processed instruction.” Found in ‘505 patent, Claim 1

The parties agree that the function is: receiving a signal indicative of a destination address of an instruction processed by the processor core and, when said instruction is of a discontinuity type, a return address of said processed instruction, and the structure is: inter-module bus 28 and address FIFO 52, and equivalents thereof.


IV. ‘036 PATENT TERMS

A. Overview of ‘036 patent

The ‘036 patent is directed to arbitrating the selection of pending interrupts when multiple interrupts occur at the same time by establishing a priority scheme. A block diagram of apparatus for handling interrupts is shown in Figure 1:

FIG. 1

A register 10 has encoded data that assigns a priority level to various interrupt requests. A register 25 has fields corresponding to multiple hardware interrupts. A mask decoder 30 is used to assign the priority of each pending interrupt in register 25. The logic in mask decoder 30 is shown in Figure 3 for one of multiple decoder circuits required for the multiple interrupts of register 25:

FIG. 3

An output of register 25 is applied by line 106 as an input to each of AND gates 103. The two bits for each interrupt set in register 10 are applied by lines 107/108 as inputs to AND gates 101, 103-2, and 103-1 as shown. Based upon the applied value of the two bits from register 10 and the presence of a pending interrupt in register 25 as signaled by line 106, the outputs 104 of AND gates 103 will have outputs corresponding to an assigned priority level for the pending interrupt.

When there are two or more pending interrupts, a selection of one interrupt for servicing is made using the leading bit detectors (LBD 1-3). The LBD outputs are applied to select circuit 60 and to arbitration circuit 50, which in turn provides enable inputs to select circuit 60. An interrupt enabled for servicing is taken at line 65.

B. Disputed terms in the '036 patent

1. "A first register having a plurality of interrupts each being programmed at one of a plurality of priority levels." Found in '036 patent, Claim 1.
2. "A second register in which a state is set for certain of said plurality of interrupts to be in a pending state." Found in '036 patent, Claim 1.

As these terms are somewhat similar, the court will deal with them together. The parties' proposed constructions for these terms are as follows:

Term	Plaintiffs' Proposals	Defendants' Proposals
“A first register . . .”	A storage area divided into a series of memory fields, each corresponding to an interrupt and programmed to indicate one of multiple priority levels.	A first register having a plurality of interrupts, the first register containing information used to disable and determine the order in which each interrupt would be serviced.
“A second register . . .”	A storage area divided into a series of memory fields, each indicating whether the corresponding interrupt in the first register is pending.	A second register in which information is contained indicating that certain interrupts are in a pending state.

Plaintiffs' position is that the terms read in view of the supporting description of the specification are clearly understood as a structure to store priority information in a series of fields (first register) and as a structure storing data indicative of whether interrupts are pending (second register).

Defendants argue that Plaintiffs' proposed construction seeks to remove the requirement that the “first register” must “have” interrupts. According to Defendants, in order to distinguish the prior art Kristick reference, the claim was amended to recite that a first register “having” a plurality of interrupts and remarks were made that as claimed there is only a single channel that carries both an interrupt and its priority level.

Defendants also contend that according to the specification priority level information is stored in the first register in order to disable an interrupt as well as to indicate its relative priority level with respect to all other interrupts in the system. Thus, Defendants say their construction of “each being programmed...” as meaning “used to disable and determine the order in which an interrupt would be serviced” is correct when the phrase is read in view of the specification.

As to their construction for the “second register,” Defendants argue that their proposal gives effect to stated requirement that a state be set for certain of the interrupts to be in a pending state, whereas Plaintiffs’ proposal, which requires only that the register have memory fields “indicating whether” an interrupt is pending, reads the “set” requirement out entirely.

a. **“First register . . .”**

Plaintiffs’ proposed construction accurately reflects that a register is a storage device having a series of fields. Further, Plaintiffs’ construction for the first register requiring that it has fields corresponding to multiple interrupt types and programmed as to the priority level of each interrupt is correct. In addition, Plaintiffs’ construction for the second register requiring that it has fields indicating whether a corresponding interrupt is pending is correct. Each aspect of Plaintiffs’ construction is supported by the specification.

Defendants fail to first read the claim language “having a plurality of interrupts” in view of the claim language. The specification describes an interrupt as a signal asserted by an external hardware device asking a CPU to perform a routine to service an external hardware device. ‘036 patent at 1:11-17. Thus, an external device can be said to have an “interrupt” but not a register. Only a recording of the occurrence of an interrupt signal, i.e., setting of a bit in a register, exists in a register. ‘036 patent at 1:39-42. Thus, one of skill would not read the claim language as actually requiring that the first register “have” an interrupt.

Defendants’ portrayal of the prosecution history regarding “having a plurality of interrupts” is misleading. During prosecution, an obviousness rejection based on Gulick and Kristick was lodged against the claim. According to the rejection, Gulick only failed to *expressly teach the selection of an interrupt for servicing*. Kristick was relied upon for its express teaching of selecting an interrupt

for servicing and combined with Gulick as an obvious modification to Gulick to provide for interrupt selection.

In remarks, the applicant stated: “As set forth in claim 1, there is a first register that corresponds to the plurality of interrupts, e.g., 20, shown in the application.” Further, applicant stated: “A second register sets certain of the interrupts to a pending state, i.e., a state in which it is to be evaluated for selection as the interrupt to be serviced.” Doc. # 196-7 at 15. As shown in Figure 1 and described in the specification, an indication of pending interrupts is provided by register 25 and not by priority register 10. Thus, even the prosecution history indicates the construction advanced by Plaintiffs and contradicts Defendants’ construction.

Further, in the remarks, Gulick was distinguished as failing to disclose a pre-assigned priority level to each interrupt. Doc. # 196-7 at 15. The amendment made to Claim 1 was made in furtherance of that distinguishing aspect over Gulick and not in regard to distinguishing Kristick. Thus, Defendants’ attachment of the claim amendment to remarks distinguishing Kristick is misplaced.

In regard to Kristick, the remarks focused on its disclosed interrupt servicing (processing) feature cited by the examiner. In Kristick, interrupt servicing is done on the basis of interrupts from a plurality of channels or ports. United States Patent No. 4,035,780 to Kristick at 1:65-66. The interrupts are applied to input registers 101 but it is level priority circuit 100-24 of Kristick that generates signals designating the channel or port having the highest priority. *Id.* at 3:56-58, 4:52-54. The applicant specifically directed remarks regarding Kristick to the arbitration circuit provided by level priority circuit 100-24 and not the input registers 101. Specifically, applicant remarked that Kristick’s arbitration circuit receives all interrupts and a request level is separately assigned. Doc. # 196-7 at 15-16.

In Kristick, the four channels (CA0 through CA3) are assigned a priority. When there are two pending interrupt requests, the level priority circuit 100-24 outputs the channel number that has been assigned the highest priority to the selector circuit 100-30. ‘780 patent to Kristick at 4:54-60. Thus, as applicant stated, each channel in Kristick has a pending interrupt and a separately assigned priority level. However, applicant did misunderstand that each channel in Kristick has four types of pending interrupt requests. *Id.* at 3:62-4:6. But what applicant remarked was that Kristick’s arbitration circuit, which includes type selector circuit 100-30, that selects and passes interrupts for servicing receives all interrupts whereas the claimed arbitration circuit only receives pending interrupts of a pre-assigned priority (i.e., pending interrupts masked to a priority level). This arrangement is shown in Figure 1, wherein only interrupts masked to a pre-assigned priority level are applied to the interrupt priority arbitration circuit 50 and decision select circuit 60 for selection and passing.

Defendants’ constructions also use the term “register” and provide no meaningful guidance to the fact finder as to what that means. Further, Defendants’ construction for “first register” adds unsupported limitations. That is, the first register according to Defendants’ construction imposes functions of disabling an interrupt and setting an order in which interrupts would be serviced, which are nowhere in the claim. While an interrupt priority level is set by it, the first register is not claimed according to a servicing order function nor is it claimed as being a mechanism to perform an interrupt disabling function. While such operation is the result imposed in the preferred embodiment by the programming of the first register with priority levels, Defendants’ inclusion of the requirement in their construction improperly imports limitations of the preferred embodiment.

The court therefore concludes that the proper construction of this term is “a storage area divided into a series of memory fields, each corresponding to an interrupt and programmed to indicate one of multiple priority levels.

b. **“A second register . . .”**

The recitation for the term “second register” indicates a functional capability and does not specify particular contents. In operation, the second register will receive pending interrupts and memory fields in the register will be “set” to indicate pending interrupt requests. When the term is read in view of the specification, Defendants’ construction is inaccurate and should be rejected. The claim language in view of the specification indicates only a capability of the “second register” as reflected in Plaintiffs’ construction.

The court will construe this term as “a storage area divided into a series of memory fields, each indicating whether the corresponding interrupt in the first register is pending.”

3. “An arbitration circuit responsive to the data in said first and second registers.”
Found in ‘036 patent, Claim 1.

Plaintiffs argue that the term should be construed as “a circuit that identifies pending interrupts from the second register and decodes the priority of interrupts from the first register.” Defendants suggest that the proper construction is “a circuit that identifies pending interrupts from the data in the second register and decodes the priority of such pending interrupts using the data in the first register.”

According to Plaintiffs, in order for the arbitration circuit to select a pending interrupt for service, it must identify pending interrupts by examining the second register contents and must decode the interrupt priorities stored in the first register. Plaintiffs criticize Defendants’ construction as limiting the necessary data to only that stored in the first and second registers.

Defendants agree to the construction proposed by Plaintiffs with the following modifications to bring the definition in line with the claim language and written description: “a circuit that

identifies pending interrupts from the data in the second register and decodes the priority of such pending interrupts using from the data in the first register.”

The claim language itself defines the functional capability of the arbitration circuit (i.e., select and pass an interrupt for servicing). The circuitry for doing so is Interrupt Priority Arbitration circuit 50 together with decision select circuit 60. What Plaintiffs’ construction advances is different functionality than specified in the claim. The identification of interrupts in the second register and the decoding of the priority of the interrupts is the function of mask decoder 30. Plaintiffs misread the claim language and their proposed construction is rejected.

Although Defendants now adopt Plaintiffs’ construction in a modified form, Defendants’ modified construction suffers the same flaws. Thus, Defendants’ modified construction is also rejected.

Defendants’ original construction applies the claim language limitation “responsive to” in a restricted fashion as meaning that only the data in the first and second registers is “sufficient” for interrupt selection and servicing. However, arbitrator 50 shown in the specification also receives an instruction Current Interrupt Priority Level (CIPL), which is described as providing the priority level of interrupts to be passed for servicing. That is, arbitrator 50 can be instructed by the application program to select and pass only interrupts having a priority level at or greater than a threshold.⁷ Accordingly, while the arbitrator 50 is “responsive” to data in the first and second registers, that data alone may or may not be sufficient to control selection of an interrupt for servicing. Thus, Plaintiffs’ construction is overly limiting in specifying the “sufficient” limitation and should be rejected. Moreover, the claim language itself specifies that a pending interrupt must be of a “predetermined one or more of said plurality of priority levels.” That language speaks directly to the CIPL instruction

to arbitrator 50. Further, the “predetermined” requirement is not embraced by the first register and the second register recitations.

Defendants’ original construction is also overly limiting in requiring that “all data” in the registers is input to the circuit. Clearly, as shown in Figure 1, not all data is provided to arbitrator 50 and select circuit 60. Only pending interrupt requests masked to a priority level are applied as inputs.

The claim language is complete in itself to set forth the functional capability that is defining of the “arbitration circuit.” The circuit is “responsive” to register data in the sense that it operates based on the programmed priority levels of the first register and the settings in the second register for pending interrupt requests. The parties’ dispute is resolved by rejecting both constructions. Both Plaintiffs’ and Defendants’ constructions are inconsistent with the specification and the claim language. As an addition note, both parties appear to consider that the term “arbitration” need not be construed as neither includes anything about it in their respective constructions.

The court therefore concludes that no construction of this term is necessary.

C. Agreed Constructions

The parties have agreed that the term “interrupt”—found in the ‘036 patent, Claim 1—should be construed to mean “A signal indicating that an event has occurred that requires attention, or that a peripheral is ready to send or receive data.”


V. ‘762 PATENT TERMS

A. Overview of ‘762 patent

The ‘762 patent is directed to load equalization to reduce power consumption in a multiprocessor chip. According to the ‘762 patent, a controller allocates tasks to the individual processors to equalize processing load between the processors on the chip. Based on the task, the

controller sets the clock frequency to as low a level as possible for proper execution of the task. Further, the controller sets the supply voltage consistent with the proper execution of the task. Thus, the approach is to reduce power consumption by running the processor only as fast as necessary to perform the assigned task based on the task and the time allotted to perform the task.

Method claim 1 is asserted. Generally, the method of controlling power consumption involves ascertaining the time allotted for carrying out an assigned task, determining the lowest frequency required to complete the task in a timely manner, and setting a supply voltage to the lowest level that insures proper operation at the determined frequency. A diagram of such operation is shown in Figure 3:



B. Disputed terms in the '762 patent

1. "Ascertaining time allotted." Found in '762 patent, Claim 1.

Plaintiffs assert that construction of this term is necessary, because the term has a plain and readily understandable meaning.

Defendants suggest that the term be construed as "determining the required completion time." They contend that the phrase needs to be construed in the context of the patent to benefit the jury. According to Defendants, the specification describes what the phrase means when it explains that

“there is a certain time when the tasks assigned to the chip must be finished” and it describes that the operating system “needs to ascertain the required completion time.”

Defendants’ construction differs from the claim language. Whatever a ”completion time” might be for a task, such amount of time is not necessarily the same as an amount of time that is “allotted” to performing the task. Defendants’ proposed construction is inconsistent with the claim language.

Defendants present their cited quotes to portions of the specification in support of their construction without providing the surrounding context in which the quotes appear. The description given as to “completion time” is in regard to the embodiment of Figure 2, which has multiple processing elements. The quote cited by Defendants is in regard to a “presumption” that there is a certain time when tasks assigned to the chip must be finished. ‘762 patent at 3:47-49. According to the specification, there might not be any particular requirement in that regard. *Id.* at 3:49-50. Further, as explained, in the illustrated embodiment of Figure 2, the expected completion time is controlling in determining a reduction in frequency. *Id.* at 3:54-65.

Although ascertaining a time allotted could be based on the required time of completion, they are not co-extensive. A time allotted should be longer than the required time of completion. If the applicant had desired the term to mean the same as “time of completion,” the applicant would have written the claim to so state.

Plaintiffs are correct that this term does not need to be construed.

2. “Assigned task.” Found in ‘762 patent, Claim 1.

Plaintiffs assert that no construction is necessary, because the term is readily understood. If the court is inclined to construe the term, Plaintiffs suggest “operation given to a sub-circuit to

perform.” Defendants propose that the term be construed as “a collection of one or more instructions allocated as a unit of work to an available processor.”

Plaintiffs dispute Defendants’ construction as rewriting the claim to substitute “processor” for “sub-circuit.” Further, a task can be achieved by execution of but a single instruction and does not require “a collection of instructions” that Defendants’ construction imposes.

Defendants argue that the term has a specific technical meaning in the context of the ‘762 patent that a jury is unlikely to understand without guidance. According to Defendants, the specification characterizes a “task” in relation to the number of instructions that need to be executed to achieve a result and as an entry on a “to-do” list, which makes it a discrete “unit of work.” Finally, Defendants argue that the specification also makes a “task” an item that is allocated to an available processor.

Clearly, a task is one or more instructions that need to be executed. See ‘763 patent at 3:13-15 (“If the number of instructions that need to be executed for each task is known...”). Also, the description implies that a set of task instructions is handled as a unit and the computational load imposed is work. Moreover, the claim characterizes an assigned task as being subject to “execution.” Thus, a “task” in the context of the ‘762 patent inherently means one or more executable instructions. Also, the plain meaning of “task” is a piece of work to be done. Finally, the plain meaning of “assigned” inherently conveys that an allocation is made. What Defendants’ construction does is add a further step of “allocating an assigned task to *an available processor*,” which improperly imports a limitation from the preferred embodiment.


The court therefore construes this term to mean “one or more instructions to be executed by the sub-circuit.”

VI. '725 PATENT TERMS

A. Overview of the '725 patent

The '725 patent is directed to accessing both stack-based storage spaces (memory)³ and register based storage spaces (memory)⁴ using a shared data path to process a set of instructions. Thus, there are two architecturally distinct storage spaces in such an arrangement. With such an arrangement, instructions for a processor to use both distinct storage spaces can be interspersed within a set of code. In the disclosed embodiment, a stack is used to store control operands and a register is used to store computational operands. The stack or the register is selected using a tag bit as part of an instruction.

An implementation of a shared data path processor is shown in Figure 4:



³ Stack-based memory is an area of memory where data is added and removed on a “last-in, first out” (LIFO) basis. An address is not required to access operands. '725 patent at 2:2-4, 2:13-15.

⁴ In register-based memory, operands are loaded from main memory and stored in a register file. An explicit address is required to access operands. '725 patent at 2:25-31.

The ‘725 patent describes that a “data path” is a set of execution units utilized to process an instruction. Further, the patent describes that a “shared data path” is a data path for which at least one element is shared between different storage spaces, such as a stack and a register file.

B. Disputed terms in the ‘725 patent

1. “Stack-based storage space.” Found in ‘725 patent, Claims 1, 18.
2. “Register-based storage space.” Found in ‘725 patent, Claims 1, 18.

The parties’ proposed constructions for these two related terms are as follows:

Term	Plaintiffs’ Proposal	Defendants’ Proposal
“Stack-based storage space.”	“Memory circuitry accessed in a last-in-first-out manner by instructions that do not specify the address of the operands on which they operate.”	“A storage structure which operates as a last-in-first-out data structure, that is a separate hardware component from the ‘register-based storage space.’”
“Register-based storage space.”	“Memory circuitry accessed by instructions that specify the address of the operands on which they operate.”	“A storage structure in which the components are addressed by their explicit number in the set, that is a separate hardware component from the ‘stack-based storage space.’”

Plaintiffs state that the parties agree the term “stack” designates a LIFO memory, but that Defendants’ construction for this term defines what a “stack” is *and* also what physical form it must take relative to the “register.” Plaintiffs also argue that Defendants’ construction for “register” imposes a limitation as to the physical form it must take relative to the “stack.”

Defendants contend that Plaintiffs should not be permitted to point to the same physical hardware to satisfy both types of storage spaces. According to Defendants, the claim language itself negates that reading by requiring two physically and structurally separate storage spaces to exist at

the same time. Defendants argue that nowhere does the specification support Plaintiffs' reading of the claim language.

The court finds Plaintiffs' constructions for both terms to be accurate. Defendants essentially raise an infringement, not a claim construction, issue. Defendants are correct that the claim requires two separate hardware limitations for literal infringement, but Plaintiffs do not dispute that point. Infringement can be shown literally by two separate hardware structures, but there need not be a one-to-one correspondence between the claimed device and the accused device. *Intel Corp. v. Int'l Trade Comm'n*, 946 F.2d 821 (Fed. Cir. 1991) ("The doctrine of equivalents does not require a one-to-one correspondence between the accused device and that disclosed in the patent."). Thus, infringement can be shown under the doctrine of equivalents when separate claim limitations are combined into a single component of the accused device. *Sun Studs, Inc. v. ATA Eqip. Leasing, Inc.*, 872 F.2d 978 (Fed. Cir. 1989), *overruled on other grounds by A.C. Aukerman Co. v. R.L. Chaises Construction Co.*, 960 F.2d 1020, 1038-39 (Fed. Cir. 1992). It is improper to include a "separate hardware component" limitation as requested by Defendants.

The court therefore construes the term "stack-based storage space" as: "memory circuitry accessed in a last-in-first-out manner by instructions that do not specify the address of the operands on which they operate." The court construes the term "register-based storage space" as: "memory circuitry accessed by instructions that specify the address of the operands on which they operate."

3. "Architecturally-distinct storage spaces." Found in '725 patent, Claims 18, 23.

Plaintiffs suggest that this term should be construed as: "memory spaces accessed by instructions from different processor architectures using different addressing techniques." Plaintiffs argue that while the parties agree what structure the term designates, there is a dispute as to what physical form the structure must take. According to Plaintiffs, different storage spaces may be

implemented using common hardware. Plaintiffs also contend that “architecture” refers to the organization and addressability of the storage space, such that two storage space architectures are distinct if their addressability is different.

Defendants propose that the court construe the term to mean “memory spaces accessed by instructions from different processor architectures and which use separate hardware components.” This is the same argument Defendants put forth for “stack-based storage space” and “register-based storage space”: that the storage spaces must be separate hardware components.

Contrary to Plaintiffs’ argument, the specification does not relate memory or storage space “architecture” to addressability. Memory or storage space “architecture” as used in the ‘725 patent describes how the processor data storage spaces are implemented as to a matter of organization. Thus, the specification refers to “register-based architecture” and “stack-based architecture.” ‘725 patent at 2:25, 2:36. 5 Plaintiffs’ construction improperly imports a limitation of “using different addressing techniques” from the preferred embodiment.

Both parties’ constructions include a requirement of being “accessed by instructions from different processor architectures.” First, the limitation is one of how the storage spaces are used and not what they are. Second, there is no dispute to be resolved that requires that limitation to be expressed. Inclusion of the limitation unnecessarily complicates the construction for the jury.

The dispute between the parties is simply whether “architecturally-distinct” requires separate hardware structures (i.e., at least two) to exist or can be provided by a single hardware structure common to both a stack-based form of memory and a register-based form of memory. Defendants say “yes,” and Plaintiffs says “no.”

Non-asserted dependent Claim 2 adds a further limitation to asserted independent Claim 1 that the storage spaces are “architecturally-distinct spaces.” This limitation is also present in asserted

independent Claims 18 and 23. This further limitation imposes a requirement that the storage spaces be physically separate hardware circuit structures within the processor, a requirement that is not imposed by the language of Claim 1, which merely recites “stack-based” and “register-based” storage spaces.

Defendants’ construction requires the storage spaces to be separate “components.” In the context of the ‘725 patent, “component” would be understood to be an “electronic component.” The understanding of “electronic component” is that it is a basic discrete device in an electronic system (e.g., active devices such as diodes and transistors and passive devices such as resistors and capacitors). The claimed “processor” in the context of the ‘725 specification is itself an electronic component and the storage spaces are on-chip memory. Thus, Defendants’ construction is overly limiting and inconsistent with the specification in specifying “components.”

As with the “stack-based” and “register-based” claim construction disputes, there is an infringement issue aspect. This is reflected by Plaintiffs’ contention that the “architecturally-distinct storage spaces” can be implemented using common hardware. As discussed with respect to the “stack-based” and “register-based” limitations of Claim 1, infringement of a common hardware realization of the claimed storage spaces can be pursued under the doctrine of equivalents. However, Claims 18 and 23, because of the “architecturally-distinct” limitation expressly set forth in those claims, presents a different situation. This limitation is not a requirement solely for literal infringement. This limitation applies under the doctrine of equivalents as well. In Claims 18 and 23, the “architecturally-distinct” limitation presents an additional limitation regarding the relationship of the storage spaces. The limitation specifically excludes storage spaces are formed in a single hardware circuit structure (i.e., the storage spaces are not physically separate hardware circuit

structures). The concept of equivalency cannot embrace a structure that is specifically excluded from the scope of the claims. *Dolly, Inc. v. Spalding & Evenflow Cos., Inc.*, 16 F.3d 394 (Fed. Cir. 1994).

Defendants' construction modified to read "physically separate hardware circuit structures" rather than "components" would be appropriate. Therefore, the court will construe this term as: "memory spaces accessed by instructions from different processor architectures and which use physically separate hardware circuit structures."

4. "A datapath which is at least partially shared by the first and second storage spaces."
 Found in '725 patent, Claim 1.
5. "A datapath at least a portion of which is selectively couplable to each of the first and second storage spaces." Found in '725 patent, Claims 18, 23.

The parties' proposed constructions for these two terms are the same. Plaintiffs suggest that the proper construction is "a set of execution units, which are utilized to process an instruction, and for which at least one execution unit is shared between different storage spaces." Plaintiffs argue that the terms are defined in the specification, and that Defendants' proposal contains a new term ("processing element") which is improper.

Defendants propose that the term be construed as "a set of execution units which are utilized to process an instruction and for which at least one processing element is shared between different storage spaces such as, for example, a stack and a register file." Defendants' position is that Plaintiffs wrongly replace the word "element" in the specification text with the phrase "execution unit." According to Defendants, the term "element" is broader than the phrase "execution unit."

Both parties deviate from the description given in the '725 patent specification at 4:51-54: a "shared" datapath "refers to a datapath for which at least one element is shared between different storage spaces such as, for example, a stack and a register file." A datapath itself is defined as a "set of execution units, which are utilized to process an instruction." *Id.* at 4:49-51.

The court therefore construes these terms to mean: “a set of execution units, which are utilized to process an instruction, and for which at least one element is shared between different storage spaces such as, for example, a stack and a register file.”

6. “Computational operands.” Found in ‘725 patent, Claim 1.

7. “Control operands.” Found in ‘725 patent, Claim 1.

The parties’ proposed constructions for these two related terms are as follows:

Term	Plaintiffs’ Proposal	Defendants’ Proposal
“Computational operands associated with one or more of the plurality of instructions.”	“Objects of a computational operation.”	“Values stored in a register file distinct from the hardware elements storing the stack.”
“Control operands associated with one or more of a plurality of instructions.”	“Objects of a control operation.”	“Values stored in a stack distinct from the register file which stores computational operands.”

Plaintiffs argue that an operand is the object of an operation and, as such, the terms are self-explanatory. Plaintiffs criticize Defendants’ constructions as importing an irrelevant limitation of completely separate physical hardware. Defendants seek the same limitation of physically separate storage structures that it sought in connection with several of the terms already construed.

There is no need to include a hardware configuration limitation in these constructions. To the extent there are separate storage structures containing each type of operand, Claim 1 otherwise sets forth that limitation. Defendants’ constructions do not seek to construe the terms, but rather to import a limitation of the preferred embodiment showing that the stack memory and the registry memory to be physically separate hardware circuit structures.

The court therefore construes the “computation operands . . .” term as: “objects of a computational operation.” The “control operands . . .” term is construed as: “objects of a control operation.”

VII. TERMS IN THE “POWER ISLAND PATENTS” (THE ‘306, ‘680, ‘885, AND ‘811 PATENTS)

A. Overview of the “power island” patents

The “power island” patents have a common specification and are directed to power saving in an integrated circuit by voltage and frequency management for discrete portions called “power islands.” A power manager determines a target power level and then determines an action to change power consumption to the target power level. Such action may be to change a frequency of operation or change supply voltage. A diagram of a system using the power saving techniques is shown in Figure 1:




FIG. 1

B. Disputed terms in the “power island” patents

1. “Selecting a frequency.” Found in ‘306 patent, Claim 16.
2. “The action including changing a frequency of operation or changing a supply voltage.” Found in ‘811 patent, Claim 1.

The parties’ proposed constructions for these two related terms are as follows:

Term	Plaintiffs’ Proposal	Defendants’ Proposal
“Selecting a frequency.”	“Selecting from a group including two or more non-zero frequencies.”	Plain and ordinary meaning, which is choosing a frequency of operation for the power island.
“The action including changing a frequency of operation or changing a supply voltage.”	“The action including changing to a frequency of operation chosen from a group including two or more non-zero frequencies or changing a supply voltage.”	Plain and ordinary meaning, which is choosing a frequency of operation for the power island.

Plaintiffs rely on the prosecution history for their construction. According to Plaintiffs, prior art “clock gating” (i.e., turning a clock on and off) was distinguished by an amendment requiring “selecting a frequency.” Plaintiffs contend that their construction captures the meaning of the term in view of the intrinsic record, whereas Defendants’ construction fails to observe the distinction over “clock gating” made during the prosecution.

Defendants argue that Plaintiffs do not construe the term but instead adds a requirement that the frequency be from a group of non-zero frequencies. Defendants contend that the prosecution history distinguishing clock gating does not support a limitation expressing a requirement for two non-zero frequencies.

Plaintiffs are seeking to give effect to its disclaimer of clock gating. However, as Defendants point out, Plaintiffs’ construction improperly adds a limitation that is not fairly based on the

disclaimer. First, a zero frequency signal would be a DC signal, which is not time-variant, and actually not a frequency signal at all. A frequency signal is necessarily time-variant in that it has a number of cycles per unit of time. Thus, “non-zero frequency” adds nothing to an understanding of the phrase. Second, clock gating does not control the frequency of a signal but instead effectively acts as an on/off switch for passing or blocking a frequency signal. The term “frequency” inherently carries a meaning that an AC and not a DC signal is involved. Thus, Plaintiffs’ reference to “non-zero frequency” adds no clarity of meaning, and Plaintiffs’ construction should be rejected.

The crux of Plaintiffs’ position is that “selecting” and “changing” do not include “clock gating” (i.e., on/off switching to pass or block a frequency signal). Defendants, however, do not suggest otherwise by expressing their construction as “choosing.”

The court therefore construes both terms according to their plain and ordinary meaning, namely “choosing a frequency of operation for the power island.”

3. “Power island(s).” Found in ‘306 patent, Claim 16; ‘680 patent, Claim 1; ‘885 patent, Claim 8; and ‘811 patent, Claims 1, 30.

Plaintiffs suggest that this term should be construed as “a section of the integrated circuit with independently controllable voltage and independently controllable frequency selected from a group including two or more non-zero frequencies.” Plaintiffs’ position is that the term means a section that is distinguished from a voltage island in power is managed control of both voltage and frequency.

Defendants propose the term be construed to mean “any section, delineation, partition, or division of an integrated circuit where power consumption is controlled within the section, partition, or division by selecting a frequency.” Defendants’ position is that the term is defined in the specification, and Plaintiffs’ construction fails to observe that definition.

Both parties' constructions largely track the definition in the specification. Plaintiffs' construction incorrectly includes reference to "non-zero frequency," which is superfluous to an understanding of the term "power island." Plaintiffs are correct, however, that both voltage control and frequency control are contemplated by the term when read in view of the intrinsic evidence.

The court therefore construes this term as: "a section of an integrated circuit where power consumption is controlled by independently varying its operating voltage and its frequency of operation."

4. "Synchronously." Found in '811 patent, Claim 30.

5. "Asynchronously." Found in '811 patent, Claim 30.

The parties' proposals for these related terms are as follows:

Term	Plaintiffs' Proposal	Defendants' Proposal
"Synchronously."	"Using a common clock."	"Using a clock."
"Asynchronously."	"Without using a common clock."	"Without a clock."

Plaintiffs argue that the terms are well known technical terms. Synchronously identifies that there is a matching of timing between components by use of a common clock. If operation of components is not synchronous, then it is termed asynchronous and their operation proceeds without a matching of timing.

Defendants dispute that there is a requirement of a common clock in order for components to operate "synchronously." Further, Defendants argue that what Plaintiffs refer to in regard to use of a common clock is "synchronized" operation of components and not components operating "synchronously."

Plaintiffs' construction is overly limiting of the terms in specifying a "common" clock. A matching of timing between components does not necessarily require a common clock. The context of use of the terms in Claim 30 is communication between power islands. In the context of communications, there must be an agreed timing scheme for a "synchronous" communication to take place. This is typically done by synchronizing the clocks in both components before a communication transmission begins through a coordination protocol. In "asynchronous" communication, there is no coordination between the components before communication transmission begins. As Defendants point out, use of a common clock is required in "synchronized" operation of two components and is not an inherent aspect of communication between two components.

Both parties erroneously focus on use of a clock or without use of a clock. In the context of communications between components, the focus is on whether coordination between components before communication transmission begins.

Therefore, the court construes "synchronously" to mean: "with coordination between components (i.e., power islands) before communication transmission begins." "Asynchronously" is construed to mean "without coordination between components (i.e., power islands) before transmission begins."

6. "Power manager." Found in '306 patent, Claim 16; '680 patent, Claim 1; '885 patent, Claim 8.
7. "Power manager to control the power consumptions." '811 patent, Claims 1, 30.

Plaintiffs contend that these terms are not technical, and are readily understood, requiring no construction.

Defendants suggest the following construction for both terms: “circuitry, device, or system configured to determine a target power level of a power island, determine an action to change the target power level, and perform an action to change the consumption power level to the target power level.” Defendants argue that the term “power manager” is expressly defined in the specification, and that the applicant therefore sought to be his own lexicographer.

Defendants are correct that the specification gives a specific definition of the “power manager.” *See, e.g.*, ‘306 patent at 4:18-28. The court will construe these terms to mean: ““circuitry, device, or system configured to determine a target power level of a power island, determine an action to change the target power level, and perform an action to change the consumption power level to the target power level.”

8. “Determine a target power level.” Found in ‘306 patent, Claim 16.

Plaintiffs contend that no construction of this term is necessary, arguing that the concepts of “power” and “target” are readily understandable.

Defendants suggest that the term should be construed as “determine a desired, calculated, or specified power consumption for a power island.” Defendants argue that the specification provides a precise definition of the term.

Defendants are correct. *See, e.g.*, ‘306 patent at 4:31-34. The court will construe this term as “determine a desired, calculated, or specified power consumption for a power island.”

9. “Register capabilities of the second power manager.” Found in ‘680 patent, Claim 1.

Plaintiffs contend that no construction of this term is necessary. Should the court opt to construe it, Plaintiffs suggest “inform the first power manager of the second power manager’s capabilities, such as what the second power manager can do and what types of commands the second

power manager can service.” Plaintiffs argue that only exemplary capabilities that are not definitional need be included in the construction.

Defendants propose the term be construed as “the second power manager informs the first power manager what the second power manager can do and what types of commands that the second power manager can service.” According to Defendants, the specification explains what is meant by one power manager registering with another power manger; namely, it informs what it can do, and what type of commands it can service. *See, e.g.*, ‘680 patent at 6:19-25.

Defendants’ construction improperly limits the term “capabilities” to those of a preferred embodiment. In another embodiment, the capabilities registered could be different. On the other hand, Plaintiffs’ proposal is confusing, by providing examples from a disclosed embodiment.

The court will therefore construe this term to mean: “inform the first power manager of the second power manager’s capabilities.”

10. “Power consumption signals indicating one or more of the power consumptions.”
Found in ‘885 patent, Claim 8.

Plaintiffs contend that no construction is necessary, because the term is clear in the context of the claim language.

Defendants propose: “a signal containing information from which the power consumed by a particular power island can be determined.” Defendants contend that there is an unstated dispute as to whether the term requires power consumption to be indicated for a particular power island, or for an entire chip. Defendants believe their proposal follows from the claim language itself, but want to have conformation of their reading of the claim expressed in a construction.

Defendants are correct in their reading of the claim language. Further, Plaintiffs do not appear to dispute that power consumption is in regard to a particular power island. Claim 8 recites a

plurality of power islands having associated power consumptions. Thus, a power consumption signal will indicate the power consumption associated with a particular power island. However, Defendants' construction is overly limiting in requiring that the signal contain information from which power consumed can be determined.

The court concludes that no construction of this term is necessary.

C. Agreed constructions

The parties have agreed to the following constructions:

1. "Independently controlled within each of the power islands." Found in '306 patent, Claim 16.

The parties agree that no construction of this term is necessary.

2. "Based on needs and operation of the integrated circuit." Found in '306 patent, Claim 16.

The parties agree that this term should be construed as: "based on future operational needs and current operating status of the integrated circuit."

3. "Each of the power consumptions adapted to be independently controlled." Found in '811 patent, Claims 1, 30.

The parties agree that no construction of this term is necessary.

4. "Back biasing." Found in '811 patent, Claim 14.

The parties agree that this term should be construed as: "biasing of the body of a device independently of the source by applying a voltage."

VIII. CONCLUSION

The above-cited claim terms of the patents-in-suit shall be construed in accordance with this Order.

SIGNED this the 29th day of April, 2013.



Keith F. Giblin
KEITH F. GIBLIN
UNITED STATES MAGISTRATE JUDGE